Ego-Surfing First-Person Videos

Ryo Yonetani, *Member, IEEE,* Kris M. Kitani, *Member, IEEE,* and Yoichi Sato, *Member, IEEE*

Abstract—We envision a future time when wearable cameras are worn by the masses, recording first-person point-of-view (POV) videos of everyday life. While these cameras can enable new assistive technologies and novel research challenges, they also raise serious privacy concerns. For example, first-person videos passively recorded by wearable cameras will necessarily include anyone who comes into the view of a camera – with or without consent. Motivated by these benefits and risks, we developed a self-search technique tailored to first-person videos. The key observation of our work is that the egocentric head motion of a target person (*i.e.*, the self) is observed both in the POV video of the target and observer. The motion correlation between the target person's video and the observer's video can then be used to identify instances of the self uniquely. We incorporate this feature into the proposed approach that computes the motion correlation over densely-sampled trajectories to search for a target in observer videos. Our approach significantly improves self-search performance over several well-known face detectors and recognizers. Furthermore, we show how our approach can enable several practical applications such as privacy filtering, target video retrieval, and social group clustering.

Index Terms—First-person video; people identification; dense trajectory.

arXiv:1606.04637v1 [cs.CV] 15 Jun 2016

1 Introduction

New technologies for image acquisition, such as wearable eye glass cameras or lapel cameras, can enable new assistive technologies and novel research challenges but may also come with latent social consequences. Around the world hundreds of millions of camera-equipped mobile phones can be used to capture special moments in life. Novel wearable camera technologies (*e.g.*, the Google Glass or the Narrative lapel camera) also offer a new paradigm for keeping a visual record of everyday life in the form of *first-person point-of-view (POV) videos* and can be used to aid productivity, such as automatic activity summarization [1], [2], [3], [4] and assistive systems [5], [6], [7].

However, as a mobile phone capturing an image with a GPS position embedded EXIF tag could be used as a means of violating one's privacy, wearable cameras also hold the inevitable risk of unintended use. Namely, first-person POV videos passively recorded by the wearable cameras will necessarily include anyone who comes into the view of a camera – with or without consent. Without proper mechanisms and technologies to preserve privacy, wearable cameras run the risk of inadvertently capturing sensitive information.

Keeping in mind both the benefits and risks of using wearable cameras, we argue that one important technology to develop is the ability to search large repositories of first-person POV videos automatically to find a target person. Much like 'ego-surfing' to search on the Internet for our own name, we believe that self-search in first-person videos can empower users to monitor and manage their own personal data. To this end, we develop a video-based self-search technique tailored to first-person videos. Because the appearance of people in first-person videos often comes under motion blur and extreme head/body pose changes (see Fig. 1), a robust approach beyond what can be accomplished by face recognition alone [8], [9], [10] is required.

Fig. 1. Robust search results. Target instances (**yellow arrows**) are detected (**unmasked regions**) despite motion blurred, non-frontal faces (**second column**) or extreme body poses (**third column**) where face detection (**green rectangles**) and recognition (**red rectangles**) fail.

To deal with the high variability of self appearance in first-person POV videos, we propose using motion as our primary feature. The key insight of our work is that a first-person video of a target user can act as a unique identifier to enable a target-specific search over a repository of first-person videos. We give a concrete example in Fig. 2. Consider a case where the target (self) individual A is conversing with another person (whom we will call observer B). When A shakes his head, it induces large global motion

- *R. Yonetani and Y. Sato are with the Institute of Industrial Science, the University of Tokyo, Japan.*
  E-mail: {yonetani, ysato}@iis.u-tokyo.ac.jp
- *K. Kitani is with the Robotics Institute, Computer Vision Group at Carnegie Mellon University, USA.*
  E-mail: kkitani@cs.cmu.edu

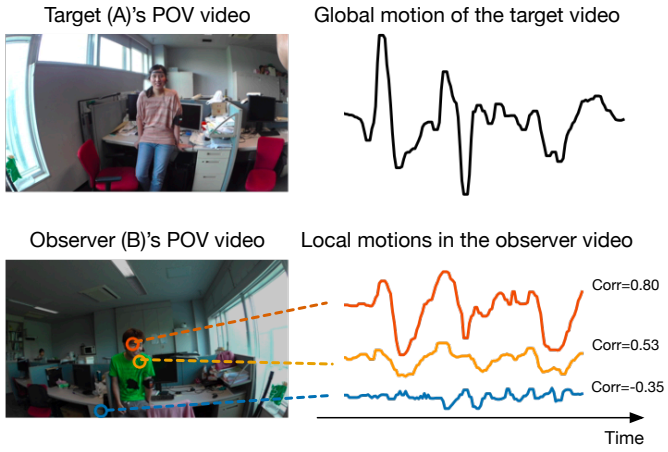

Fig. 2. Global motion in the target (A)'s points-of-view video and local motions in the observer (B)'s video. The local motion in a target region (**red** line plot) has a higher correlation with the global motion (**black** line plot) compared with that in non-target regions (**yellow** and **blue** plots).

(the camera moves according to the black line plot in Fig. 2) in the video. Now, from the perspective of observer B, we expect to see the same shake pattern but in the form of a local motion pattern in B's POV video (we see the target individual A shaking his head at the red circle in Fig. 2). This correlation between the global motion of target A's video and the local motion of observer B's video indicates that these two videos are indeed related. Furthermore, this correlation is expected to increase only in target regions. This illustrates the general insight that the ego-motion of a target is a unique signature that can be used to localize the self in observer videos.

Based on this insight, we develop a novel motion correlation-based approach to search and localize a target individual in a collection of first-person videos. The algorithm takes as input the target's POV video and retrieves as output all instances of that target individual from an observer's POV videos. Our algorithm proceeds as follows. First, densely-sampled point trajectories are generated over all videos as target candidates. Second, each trajectory is evaluated to compute their 'targetness' based on the correlation between the local motion pattern along the trajectory and the global motion pattern of the target video. Third, a supervised classifier on trajectories refines the targetness scores by taking into account the generic targetness of the trajectories. The targetness evaluation is posed as a binary-class Bayesian inference problem where the likelihood and prior are given by the motion correlation and trajectory classifier, respectively.

Experimental results show that our approach significantly improves the self-search performance over baseline face recognizers. Furthermore, to demonstrate the potential impact that self-search can have on assistive technologies, we applied our method to three tasks: (1) privacy filtering, (2) target video retrieval, and (3) social group clustering.

Related Work

The idea of searching for a specific person in images or videos has been addressed in several areas of computer vision. One example is *person re-identification* in the context of visual surveillance. An extensive survey of the field can be found in [11], [12], [13]. One common approach for person re-identification is to utilize visual signatures of a specific person, such as color and texture, to find specific individuals in images. Because many approaches presuppose a surveillance scenario, the features and approaches are often adapted for videos captured by a static camera (single POV, constant background, *etc.*). With the exception of work using active cameras [14], re-identification approaches are not designed to deal with extreme camera motion.

Work in the area of *egocentric vision* has utilized person identification as a feature for analyzing human interactions. Many studies have relied on off-the-shelf face detectors and recognizers [2], [15], [16], [17], [18], [19], [20], [21], [22]. In many of these scenarios, the use of face detection is justified since people are engaged in conversation and a wearable camera is relatively stable.

Poleg *et al.* [18] has proposed a person identification method based on the correlation of head motion in first-person videos. Their method relies on people detection to track the head of each target candidate, making it challenging to perform the identification task reliably when a person is mobile and when significant ego-motion is present. By contrast, our approach does not require people detection but directly examines the correlation over densely-sampled point trajectories. This makes our method perform well even with videos captured under significant head motion.

Another approach for identifying specific individuals is the use of geometric information in 2D [23], [24] or 3D [25], [26], [27]. An accurate geometric map can be used to compute the precise location of individuals with wearable cameras, and the location can be used to estimate visibility in other cameras. However, these approaches often require preliminary scanning of scenes (*e.g.*, [25]), making it applicable only in pre-recorded and static places.

Contributions

In this work, we extend our prior work presented in [28]. To the best of our knowledge, [28] is the first to address the topic of self-search in first-person videos with significant ego-motion, where face recognition and geometric localization are not applicable (*i.e.*, people with high facial appearance variability, captured by the cameras with significant motion, without any restriction on recorded places). In addition, this work makes the following contributions:

- Unlike our previous approach [28], we introduce a new form of target candidates based on densely-sampled point trajectories in Sec. 2. Compared with supervoxel-based candidates used in [28], trajectory-based candidates are comparably robust against the variability of faces, and they can be generated much faster. Experiments in Sec. 4 demonstrate the effectiveness of our new approach in terms of both accuracy and efficiency.
- We propose a two-step search scheme to accelerate targetness evaluation in Sec. 3. An upper-bound of correlation-based targetness is rapidly estimated in the first step to limit the number of candidates to apply a stable but slow correlation evaluation in the

latter step. This makes the overall procedure more efficient while keeping high search performance.
- In addition to the task of target localization, we quantitatively evaluated the performance of our approach on the tasks of target video retrieval (searching a repository of first-person videos for the videos including target people) and social group clustering (finding a group of people interacting with each other from the repository) in Sec. 4.

2 Correlation-based Target Search

Assume that we are given a repository of first-person POV videos of people's interactions in various places. Our goal is to search for a target person (*i.e.*, the self) across the repository. More specifically, we wish to evaluate how likely each video includes the target person and to localize all the instances of the target in the video. In what follows, we consider two types of videos: videos recorded (1) by the target (*target video*) and (2) by observers (*observer videos*).

To perform searches, the target video is used as a unique identifier to compare against all the observer videos. Candidates of target instances are generated from observer videos. Global motion induced by ego-motion in the target video is matched against local motion observed in the candidates. Matched candidates with high correlation are suggested as target instances.

2.1 Generating Target Candidates

Generating good target candidates is essential to achieve robust and efficient searches. Poleg *et al.* [18] generate target candidates via people detection. However, people detection often becomes difficult when heavy occlusion and motion blur are present on faces or when the faces are extremely small in observer videos. Our prior work [28] instead used a supervoxel hierarchy [29], [30]. While this approach has enabled robust search against the variability of faces, hierarchical supervoxel segmentation requires rather long computation time making it difficult to apply on a large-scale repository of first-person videos.

As a more efficient approach, a new target candidate based on densely-sampled point trajectories is proposed in this paper. That is, feature points detected in a certain frame are tracked over time to serve as the candidates. Target candidates generated in this way allow us to avoid evaluation in obvious background regions (*e.g.*, textureless walls and floors), making the overall search procedure more efficient than the supervoxel-based approach used in [28]. They are also comparably robust against the facial appearance variability because they rely on only the motion estimation.

Following [13], feature points are sampled with a uniform step e_W and selected with a threshold based on the good-feature-to-track criterion [31]. Dense optical flow fields such as [32] are then estimated to track the points. While [13] limits the length of trajectories to ensure that is short enough to avoid drifting, we want trajectories that are as long as possible to compute a motion correlation stably, as suggested in [18]. Therefore, we instead set the minimum and maximum lengths to the trajectories, $[L_{\min}, L_{\max}]$, and continue tracking as long as the trajectories are shorter than $L_{\max}$ and the tail of each trajectory satisfies the good-feature-to-track criterion. Trajectories shorter than $L_{\min}$ are just unused for target candidates. More implementation details are given in Appendix A.

2.2 Evaluating Targetness

Given a target video as a search query, we evaluate the targetness of candidate trajectories to generate a pixel-level targetness map of observer videos. In what follows, we first introduce a mathematical formulation of the targetness evaluation for each trajectory.

Let us denote a set of candidate trajectories by $\mathcal{X} = \{X^{(1)}, \ldots, X^{(N)}\}$. Each candidate trajectory $X^{(i)}$ is defined by a sequence of points $X^{(i)} = (\boldsymbol{x}_1^{(i)}, \ldots, \boldsymbol{x}_{l^{(i)}}^{(i)})$, where $\boldsymbol{x}_t \in \mathbb{R}_+^2$ is a pixel at the t-th frame, starting at the $b^{(i)}$-th frame and lasting for $l^{(i)}$ frames.

For each candidate $X^{(i)}$, we introduce a binary assignment variable $a_{X^{(i)}} \in \{1, 0\}$ that takes $a_{X^{(i)}} = 1$ if $X^{(i)}$ is in a target region. Then, the targetness of $X^{(i)}$ given a target video $\mathcal{V}_{\mathrm{G}}$ is defined by a posterior probability $P(a_{X^{(i)}} \mid \mathcal{V}_{\mathrm{G}})$. We further decompose this posterior using Bayes' rule to obtain the following Eq. (1).

$$P(a_{X^{(i)}} \mid \mathcal{V}_{\mathrm{G}}) \propto P(\mathcal{V}_{\mathrm{G}} \mid a_{X^{(i)}}) P(a_{X^{(i)}}). \tag{1}$$

We will estimate the likelihood $P(\mathcal{V}_{\mathrm{G}} \mid a_{X^{(i)}})$ using a cross correlation score (Sec. 2.2.1). The prior term $P(a_{X^{(i)}})$ refines the correlation-based targetness evaluation by learning various statistics of candidate trajectories indicating a trait of generic targetness (Sec. 2.2.2).

2.2.1 Correlation-based targetness

Each target candidate evaluates its targetness based on how much its local motion is correlated with the referential global motion of a target video. We follow [18] and evaluate the zero-mean normalized cross correlation (ZNCC) independently along horizontal and vertical directions to average them. Importantly, we show in Sec. 3 that this correlation-based targetness can be efficiently evaluated by estimating its upper-bound.

Motion estimation for target and observer videos proceeds as follows. In the target video, we first compute sparse optical flows by [33]. We follow [13] and assume that two consecutive frames are related by a homography. Then, the homography estimated from the flows is used to compute a global motion vector for every pixel. Finally, we average the global motion vectors for each frame to obtain a global motion pattern as a sequence of two dimensional (*i.e.*, horizontal and vertical) vectors. As for the observer videos, dense optical flow fields are estimated by [32] in addition to the per-pixel global motion vectors. A local motion vector for each pixel is computed by subtracting the global motion vector from the optical flow. Finally, a local motion pattern on a candidate trajectory is obtained by referring to the local motion at each point of the trajectory. To cope with motion estimation error, we apply a median filter on global and local motion patterns as post-processing.

Now, denote the local motion pattern of a candidate trajectory $X^{(i)}$ as a sequence of two dimensional vectors, $\boldsymbol{m}^{(i)} = \left((u_1^{(i)}, v_1^{(i)}), \ldots, (u_{l^{(i)}}^{(i)}, v_{l^{(i)}}^{(i)})\right)$, where $u_t^{(i)}$ and $v_t^{(i)}$

are horizontal and vertical elements of motion, respectively. Likewise, we consider the global motion pattern $\boldsymbol{M} = ((U_1, V_1), \ldots, (U_{l^{(i)}}, V_{l^{(i)}}))$ adaptively cropped according to the interval $[b^{(i)}, b^{(i)}+l^{(i)}]$ (U_t and V_t are also horizontal and vertical elements, respectively). Note that horizontal and vertical elements of $\boldsymbol{m}^{(i)}$ and $\boldsymbol{M}$ are independently normalized to have zero mean and unit variance beforehand, and the vertical elements of global motion, V_t, are inverted as they are inverted to $v_t^{(i)}$ (e.g., head motion by nodding down appear as upper global motion in a target video). Then, the correlation between $\boldsymbol{m}^{(i)}$ and $\boldsymbol{M}$ is defined by averaging horizontal and vertical ZNCCs as follows:

$$C(\boldsymbol{m}^{(i)}, \boldsymbol{M}) = \frac{1}{2}\left(\frac{\sum_t u_t^{(i)} U_t}{l^{(i)}} + \frac{\sum_t v_t^{(i)} V_t}{l^{(i)}}\right). \quad (2)$$

Finally, we scale $C(\boldsymbol{m}^{(i)}, \boldsymbol{M})$ into the range of $[0, 1]$ to obtain the likelihood defined in Eq. (1). Following [28], we use a sigmoid function for scaling; namely, $P(\mathcal{V}_{\mathrm{G}} \mid a_{X^{(i)}}) \triangleq (1 + \exp(-C(\boldsymbol{m}^{(i)}, \boldsymbol{M})))^{-1}$.

2.2.2 Data-driven generic targetness

The prior $P(a_{X^{(i)}})$ in Eq. (1) is introduced to consider generic targetness of a candidate trajectory $X^{(i)}$. For example, trajectories that track a skin-color point are more likely to be in a facial region, and stationary trajectories often come from a background region. This prior aims to learn such traits from many pairs of observer videos and corresponding masks annotating target regions. While manual annotations of target people are required here, this prior is independent of specific people and backgrounds. Therefore, the learning of the prior needs to be carried out only once and is not necessary for each target person.

Specifically, we first extract positive and negative feature samples from target and non-target regions, respectively, and then train a binary classifier. This classifier produces a posterior probability of target candidates belonging to the target class, *i.e.*, $P(a_{X^{(i)}} = 1)$. To capture a trait of generic targetness, we extract color (in the HSV space) and local motion for each pixel on candidate trajectories. The mean and standard deviation of the color and motion are calculated to serve as features. We also use the temporal length of trajectories as a feature.

2.3 Pixel-level Targetness Map

Estimating the targetness at the pixel level is important to localize target instances in observer videos. Unlike supervoxel-based candidates used in [28], trajectory-based candidates do not cover an entire video frame. Therefore, some interpolation techniques are needed to generate a pixel-level targetness map. We here introduce a simple nearest-neighbor approach to this problem. Intuitively, a certain pixel will take the targetness of its nearest neighbor candidate trajectory if the candidate is sufficiently close to the pixel.

Like a binary assignment $a_{X^{(i)}}$ defined for each candidate $X^{(i)}$, we consider a binary assignment variable $a_{\boldsymbol{x}_t}$ for each pixel $\boldsymbol{x}_t$. That is, $a_{\boldsymbol{x}_t} = 1$ if $\boldsymbol{x}_t$ corresponds to target regions in an observer video. The targetness for pixel $\boldsymbol{x}_t$ is defined by the posterior probability of $a_{\boldsymbol{x}_t} = 1$ given a target video $\mathcal{V}_{\mathrm{G}}$. The nearest-neighbor rule to interpolate $a_{\boldsymbol{x}_t}$ from $a_{X^{(i)}}$ is defined as follows:

$$P(a_{\boldsymbol{x}_t} | \mathcal{V}_{\mathrm{G}}) \triangleq \begin{cases} P(a_{X^{(\hat{i})}} \mid \mathcal{V}_{\mathrm{G}}) & \text{if } D(\boldsymbol{x}_t, X^{(\hat{i})}) \le r, \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where $\hat{i} = \arg\min_i D(\boldsymbol{x}_t, X^{(i)})$ is given by the spatial Euclidean distance:

$$D(\boldsymbol{x}_t, X^{(i)}) \triangleq \begin{cases} \|\boldsymbol{x}_t - \boldsymbol{x}^{(i)}_{t-b^{(i)}+1}\|_2 & \text{if } b^{(i)} \le t < b^{(i)} + l^{(i)}, \\ \infty & \text{otherwise.} \end{cases} \quad (4)$$

The radius threshold r is defined to be twice as large as the trajectory sampling step e_W to cover pixels sufficiently between candidates. We will use this per-pixel targetness map for the task of target localization in Sec. 4.2.

2.4 Affinity between First-Person Videos

In addition to the target localization, another important ability in target search is to retrieve videos that include a target person from a collection of first-person videos. As we assume in this work that first-person videos are recorded during interaction, such videos are highly likely to be taken by the people involved in the same interaction group. To take this into account, we propose an affinity between two first-person videos that describes how likely two camera wearers interact with each other for target video retrieval.

First, we define how likely a person p recording the video $\mathcal{V}_p$ is to interact with a person q recording the video $\mathcal{V}_q$ by $A(\mathcal{V}_q \mid \mathcal{V}_p)$. We expect that, the longer p appears in the video $\mathcal{V}_q$, the more likely p is interacting with q. Given a set of candidates $\mathcal{X} = \{X^{(1)}, \ldots, X^{(N)}\}$ generated from $\mathcal{V}_q$, $A(\mathcal{V}_q \mid \mathcal{V}_p)$ is defined as follows:

$$A(\mathcal{V}_q \mid \mathcal{V}_p) = \max_i P(\mathcal{V}_p \mid a_{X^{(i)}}) \cdot S(l^{(i)}; \mu_l), \quad (5)$$

where $S(l^{(i)}; \mu_l) = \left(1 + \exp(-l^{(i)} + \mu_l)\right)^{-1}$ is a sigmoid function with the center at μ_l to weigh each candidate based on its length. We also point out that correlation-based targetness $P(\mathcal{V}_p \mid a_{X^{(i)}})$ is more necessary to compute the affinity than data-driven generic targetness $P(a_{X^{(i)}})$ because $P(a_{X^{(i)}})$ encourages *any* candidates equally regardless of $\mathcal{V}_p$, making affinities more similar to each other.

Another important observation is that, when a target person is interacting with an observer, the observer can also be present in the first-person video recorded by the target. Namely, at least one of two people is expected to be seen in the other video if they are interacting with each other. Therefore, we define the affinity between two videos, $A(\mathcal{V}_p, \mathcal{V}_q)$, as follows:

$$A(\mathcal{V}_p, \mathcal{V}_q) = \max\left(A(\mathcal{V}_p \mid \mathcal{V}_q), A(\mathcal{V}_q \mid \mathcal{V}_p)\right). \quad (6)$$

This symmetric affinity can indeed work better than the asymmetric targetness defined in Eq. (5), as will be demonstrated in experiments. We also show in the experiments that this affinity can be further applied to cluster groups of interactions from a collection of videos.

3 Efficient Target Search

In this section, we consider accelerating target search as a main extension over [28]. Among the overall pipeline of our approach consisting of candidate generation and targetness evaluation, the latter is particularly critical because it needs to be performed against all videos in a repository as many times as a new query video is made. In what follows, we introduce an efficient evaluation scheme of targetness by estimating its upper-bound.

3.1 Two-step Evaluation of Targetness

As can be seen in Eq. (2) and Eq. (5), the time complexity of evaluating targetness is linearly proportional to the temporal length of target candidates (*i.e.*, $\mathcal{O}(l^{(i)})$ for each candidate). Namely, the longer time we need to track candidates to stably calculate a correlation, the longer time we need to evaluate their targetness.

To solve this problem, our proposal is to limit the number of candidates to perform a stable but slow correlation evaluation. More specifically, we introduce the piecewise constant approximation of motion patterns to estimate an upper-bound of correlation scores between the patterns. This upper-bound estimation for each candidate can be done in constant time regardless of the candidate length, making it possible to avoid immediately many unnecessarily evaluations of candidates resulting in a low correlation score.

Our scheme is summarized in the following two steps: (1) approximating motion patterns into K constant pieces to estimate an upper-bound of correlation scores and (2) picking up the candidates having the top P-percentile high upper-bound score to compute their actual correlations. Now, the time complexity improves to $\mathcal{O}(K)$ for each candidate and $\mathcal{O}(l^{(i)})$ for the only P percentage of all the candidates (usually $K \ll l^{(i)}$). Note that the scheme is inspired by the GEMINI (generic multimedia indexing) framework [34], [35] commonly used in the field of time-series data mining.

3.2 Estimating an Upper-bound of Correlation

Let us first introduce the definition of piecewise constant approximation. Let $\boldsymbol{u} = (u_1, \ldots, u_l)$ and $\boldsymbol{U} = (U_1, \ldots, U_l)$ be horizontal patterns of local and global motion, respectively (we omit the index of local motion patterns without loss of generality). We assume that $\boldsymbol{u}$ and $\boldsymbol{U}$ are normalized to have zero mean and unit variance. K-piecewise constant approximation of the local motion pattern $\boldsymbol{u}$ results in a K-dimensional pattern, $\bar{\boldsymbol{u}} = (\bar{u}_1, \ldots, \bar{u}_K)$, where $\bar{u}_k$ is defined as follows:

$$\bar{u}_k = \frac{K}{l} \sum_{t=(k-1)\frac{l}{K}+1}^{k\frac{l}{K}} u_t. \tag{7}$$

Note that K is set to be a divisor of l and is constant regardless of the length of original patterns[1]. In the same way, we can also obtain the approximation for the global motion pattern $\boldsymbol{U}$ as $\bar{\boldsymbol{U}} = (\bar{U}_1, \ldots, \bar{U}_K)$.

As proved in [35], the Euclidean distance between two piecewise constants provides a lower-bound of their actual Euclidean distance. Namely, the following inequality holds:

$$\frac{1}{l}\sum_t (u_t - U_t)^2 \geq \frac{1}{K}\sum_k (\bar{u}_k - \bar{U}_K)^2, \tag{8}$$

where the equality holds for $K = l$. Therefore, an upper-bound of ZNCC between $\boldsymbol{u}$ and $\boldsymbol{U}$ can be derived as follows:

$$\frac{1}{l}\sum_t u_t U_T \leq \frac{1}{K}\sum_k \bar{u}_k \bar{U}_k + \left(1 - \frac{\sigma^2_{\bar{\boldsymbol{u}}} + \sigma^2_{\bar{\boldsymbol{U}}}}{2}\right), \tag{9}$$

where $\sigma^2_{\bar{\boldsymbol{u}}}$ and $\sigma^2_{\bar{\boldsymbol{U}}}$ are the variance of piecewise constants (see Appendix B for more details). By letting $\bar{\boldsymbol{v}}$ and $\bar{\boldsymbol{V}}$ be the piecewise constant approximation of vertical patterns of local and global motion, and by leting $\sigma^2_{\bar{\boldsymbol{v}}}$ and $\sigma^2_{\bar{\boldsymbol{V}}}$ be their variances, the upper-bound of Eq. (2) can be derived as follows:

$$C(m^{(i)}, M) \leq \frac{1}{2}\left(\frac{\sum_k \bar{u}^{(i)}_k \bar{U}_k}{K} + \frac{\sum_k \bar{v}^{(i)}_k \bar{V}_k}{K}\right) + Z \tag{10}$$

$$Z = 1 - \frac{\sigma^2_{\bar{\boldsymbol{u}}} + \sigma^2_{\bar{\boldsymbol{U}}} + \sigma^2_{\bar{\boldsymbol{v}}} + \sigma^2_{\bar{\boldsymbol{V}}}}{4}. \tag{11}$$

In this upper-bound, Z increases up to 1 as l becomes larger than K. This property works well because the top P-percentile target candidates preferentially include the ones seen by observers for a longer time, and they are more likely to be actual target instances in practice. In the following experiments, we demonstrate that the proposed two-step scheme successfully reduces computation time while providing comparable localization performance to where the actual correlation is computed for all the candidates.

4 Experiments

We investigate the performance of the proposed approach on our new dataset introduced in [28][2] and CMU-group first-person video dataset used in [1], [25], [26]. The implementation details are described in Appendix A, and we would like to particularly note that our algorithm was able to run on videos of size 320x180 while baseline face detectors and recognizers required full resolution videos (1920x1080 for out dataset and 1280x960 for the CMU dataset).

4.1 Datasets

Our dataset was collected in eight different interaction scenes. The number of participants equipped with a camera was two or three for each scene. These participants stayed at the same position but often changed their poses. Interactions were recorded at 60 fps, 30 sec in four indoor (**Indoor 0** – **Indoor 3**) and four outdoor scenes (**Outdoor 0** – **Outdoor 3**). On the CMU dataset, 11 participants formed groups to play pool or table tennis or to sit on couches to chat or to talk with each other at a table. They changed their poses and positions, and often disappeared from observer videos, standing for a more challenging scenario than our dataset.

1. In practice, we trim the last several frames of each trajectory so that l is a multiple of K.

2. Our dataset and codes are available on the project page http://yonetaniryo.github.io/corrsearch/.

For the nine videos available for analysis, we used 3861st – 5300th frames (30 sec at 48 fps), where **Pool**: two people played pool, **Tennis**: three played table tennis, and **Chat**: the remaining four chatted on a couch. In total, 29 videos comprising 11 different scenes were used (see also Fig. 3).

We manually annotated image regions corresponding to a target person's head every 0.5 second. This is because the local motion corresponding to ego-motion should be observed in the region of the head. These annotations also served as a supervised label to learn the prior $P(a_{X^{(i)}})$. We used the linear discriminant analysis by following [28] because it performed the best, but any other classifier could work as well. Because we evaluated two completely different datasets, we used one for training the prior to test the other. It ensured that people and scenes in test subsets did not appear in training ones.

TABLE 2
AUC scores averaged over scenes on the CMU dataset. **ZR**: Mixtures-of-tree based facial landmark detector [36]. **VJ**: Haar-cascade face detector [37]. **VJ + LBPH**: face recognition using local binary pattern histograms [38]. **VJ + FisherFace**: face recognition using the Fisher face [39]. No face recognition results were provided for Pool data because each participant was observed in only the single session.

	Pool (2 people)	Tennis (3 people)	Chat (4 people)	All
Proposed (C+G)	**0.88**	**0.85**	0.84	**0.85**
Proposed-25	0.83	0.78	**0.85**	0.82
Proposed (C)	0.84	0.79	0.80	0.80
YKS [28]	0.83	0.80	0.79	0.80
ZR [36]	0.49	0.50	0.54	0.53
VJ [37]	0.52	0.56	0.58	0.56
VJ + LBPH [38]	-	0.52	0.53	0.52
VJ + FisherFace [39]	-	0.51	0.53	0.52

4.2 Evaluation on Target Localization

We first adopted the task of target localization. Namely, given a pair of target and observer videos (in what follows, we refer to this pair as *session*), we evaluated how much per-pixel targetness maps in Eq. (3) were correlated with ground truth masks. Among the 52 sessions in total, we did not use four sessions (two in **Outdoor 2** and another two in **Chat**), where the target and observer sat down side by side and hardly looked at each other. For the remaining sessions, we calculated the area under the receiver-operator characteristic curve (AUC) based on the pixel-wise comparisons between per-pixel targetness scores and ground-truth annotations.

In this experiment, we evaluated the following four variants of the proposed approach.

- **Proposed (C+G)**: the proposed method with a combination of correlation-based and generic targetness, without the two-step targetness evaluation in Sec. 3.
- **Proposed-25**: the proposed method with a combination of correlation-based and generic targetness, with the two-step targetness evaluation. The top P=25 percent of candidates were chosen to evaluate the actual correlations.
- **Proposed (C)**: the proposed method using only the correlation-based targetness (*i.e.*, $P(a_{X^{(i)}} = 1) = 1$ for all candidates).
- **YKS**: the method proposed in our prior work [28]. Instead of densely-sampled point trajectories, supervoxel hierarchies were used for target candidates.

In addition, several off-the-shelf face detectors and recognizers served as a baseline. We used the mixtures-of-tree based facial landmark detector [36] (**ZR**) and the Haar-cascade face detector [37] (**VJ**). **VJ** combined frontal and profile face models to permit head pose variations. We also utilized face recognizers based on local binary pattern histograms [38] (**VJ + LBPH**) and the Fisher face [39] (**VJ + FisherFace**). These two recognizers learned target faces from different sessions and ran on the detection results of **VJ**.

Results

Fig. 4 and 5 depict some results of target localization. In the first and second rows of Fig. 4, the proposed methods successfully located a target person (the woman in the pink shirt) in two different points-of-view (green circles in the figure). In the third and fourth rows, the target person (the man in the green shirt) was found regardless of the lighting conditions, while none of the two face recognizers could recognize both cases. In the bottom two rows, the proposed methods were able to distinguish between two people while the face detector detected only one. The target candidates in the background regions sometimes obtained the highest targetness score using our prior method **YKS** [28] (the fourth and sixth examples in Fig. 4 and all the examples in Fig. 5). Because our trajectory-based approaches did not evaluate obvious background regions, **Proposed (C+G)**, **Proposed-25** and **Proposed (C)** were able to work well even in such cases. The comparison between **Proposed (C+G)** and **Proposed (C)** illustrates the necessity of generic targetness. Correlation-based targetness sometimes became high in textured background regions such as in the second row of Fig. 4.

Tab. 1 and 2 describe AUC scores for each scene. Our approaches obviously outperformed baseline methods. While face recognizers eliminated incorrect detections of faces, they also failed to recognize target faces due to the high variability of facial appearances. On average, **Proposed (C+G)** improved the AUC scores over **YKS**.

On the two-step targetness evaluation

Proposed-25 performed comparably well with **Proposed (C+G)**. It was also better than **YKS** [28] and the baseline methods. To analyze its performance in more detail, we investigate how AUC scores and computation times change for different settings of P (the percentage of candidates to evaluate actual correlation-based targetness) and K (the number of pieces for the piecewise constant approximation). More specifically, we tested $P = 1, 5, 10, 25, 50, 75, 100$ and $K = 4, 8, 16, 32, 64$. Note that $P = 100$ was equivalent to **Proposed (C+G)** in that all candidates were evaluated.

First, we compared AUC scores on several different combinations of P and K in Fig. 6. The proper setting of K depended on the frame rate to record first-person videos as well as the maximum length of candidate trajectories $L_{\max}$. Because we set $L_{\max} = 1024$ at 60 fps, the piecewise constant approximation with $K = 4$ could approximate about four-second motion into one constant value, making it difficult to observe quick head motion. Nevertheless, our

Fig. 3. Representative frames of interaction scenes in our dataset and CMU dataset.

method maintained comparable AUC scores as long as P was large enough. The performance drastically dropped when $P < 25$, and became comparable with baselines when $P = 1$. When one needs to localize all the instances in observer videos, P should not be extremely small because the candidates may be incorrectly eliminated. For example, we found that the combination of $P = 25$ and $K = 64$ mostly worked fine in **Proposed-25** for a variety of scenes comprised in the two datasets.

Next, we compared computation times[3] on various P, while fixing K to $K = 64$. Within the overall procedure, generating queries and candidates requires a constant time regardless of P (see the left of Fig. 7). Importantly, even if $P = 100$, the proposed method already improved computational efficiency over our prior approach **YKS** [28] because the hierarchical supervoxel segmentation needed more than two seconds for each frame[4].

Focusing on the targetness evaluation, **Proposed-25** reduced computation times by 29 percent compared with **Proposed (C+G)** on the per-pixel evaluation. Note that accelerating this evaluation procedure was critical because it needed to be applied for each combination of target and observer videos ($N \times M$ times for N target and M observer videos), while the generation procedure needed to be done only once for each video (N or M times for the same setting). We also found a 63 percent reduction in computation times by **Proposed-25** on a per-candidate evaluation used in Eq. (6) (the right of Fig. 7). This acceleration is particularly beneficial for the tasks of target video retrieval and social group clustering in Sec. 4.3, which requires only per-candidate targetness to compute an affinity between two first-person videos.

3. We implemented an overall procedure in Python and tested it on a single thread of MacPro with a 2.7 GHz 12-Core Intel Xeon E5.

4. We used LIBSVX with the default set of parameters provided with the code (http://www.cse.buffalo.edu/ jcorso/r/supervoxels/). See [28] for more implementation details.

TABLE 1
AUC scores averaged over scenes. **ZR**: Mixtures-of-tree based facial landmark detector [36]. **VJ**: Haar-cascade face detector [37]. **VJ + LBPH**: face recognition using local binary pattern histograms [38]. **VJ + FisherFace**: face recognition using the Fisher face [39].

	Indoor 0 (2 people)	Indoor 1 (2 people)	Indoor 2 (3 people)	Indoor 3 (3 people)	Outdoor 0 (2 people)	Outdoor 1 (2 people)	Outdoor 2 (3 people)	Outdoor 3 (3 people)	All
Proposed (C+G)	**0.89**	0.91	**0.90**	**0.87**	0.86	0.82	0.85	**0.77**	**0.85**
Proposed-25	0.89	**0.93**	0.86	0.81	0.86	0.82	0.82	0.79	0.83
Proposed (C)	0.88	**0.92**	0.85	0.81	0.85	0.76	0.82	0.73	0.81
YKS [28]	0.88	0.89	0.75	0.72	**0.91**	**0.88**	**0.90**	0.70	0.79
ZR [36]	0.57	0.60	0.60	0.64	0.63	0.66	0.69	0.62	0.62
VJ [37]	0.66	0.73	0.66	0.62	0.75	0.77	0.73	0.62	0.67
VJ + LBPH [38]	0.50	0.50	0.51	0.54	0.57	0.55	0.55	0.50	0.52
VJ + FisherFace [39]	0.50	0.50	0.51	0.54	0.50	0.52	0.50	0.51	0.51

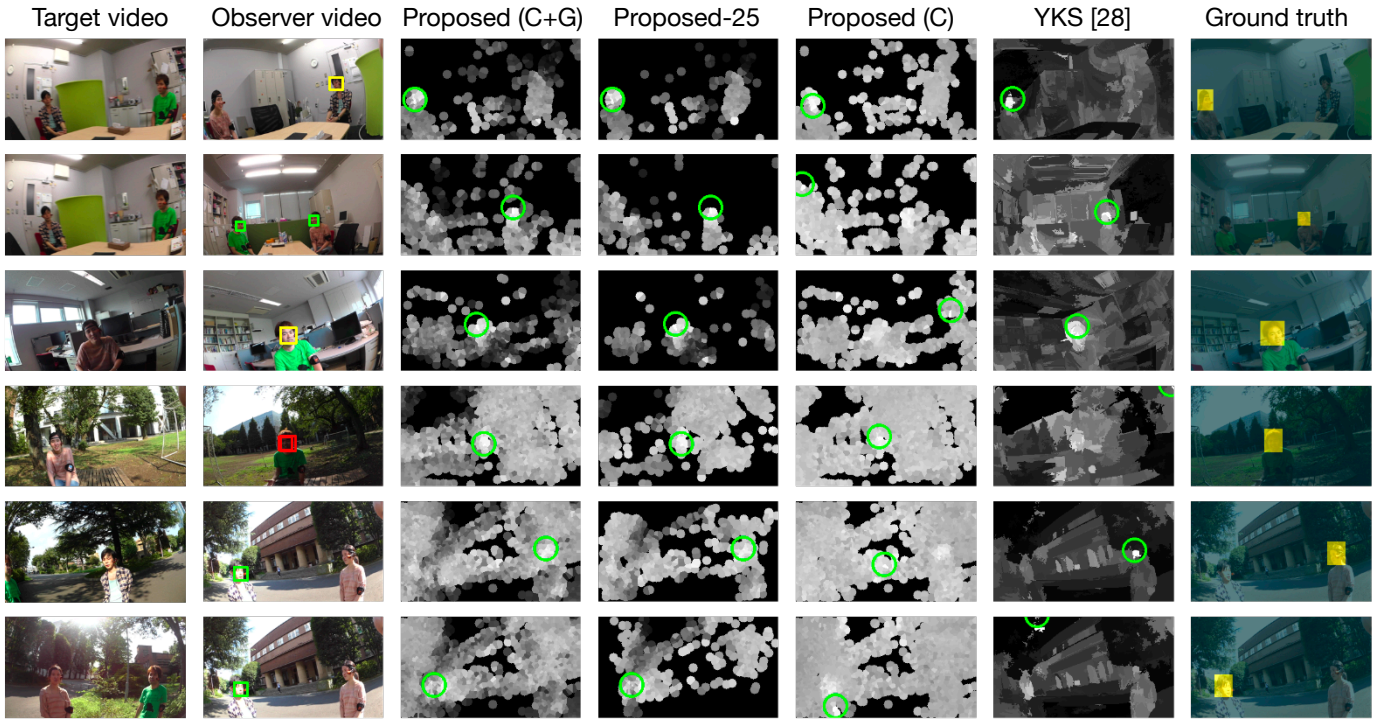

Fig. 4. Target localization results of our dataset. The target candidate with the highest targetness score is specified by the green circle on **Proposed (C+G)**, **Proposed-25**, **Proposed (C)** and **YKS** [28]. **ZR**: Mixtures-of-tree based facial landmark detector [36]. **VJ**: Haar-cascade face detector [37] (green rectangles in the second column). **VJ + LBPH**: face recognition using local binary pattern histograms [38] (red rectangles). **VJ + FisherFace**: face recognition using the Fisher face [39] (yellow rectangles).

Target segmentation for privacy filtering

Target localization results can be used to segment target regions in videos. We ran GrabCut [40] on per-pixel targetness maps and found the region with the highest targetness score. Such segmentation is applicable for privacy filtering. For example, hiding regions other than target ones such as in Fig. 8 can prevent us from violating the privacy of other people who accidentally come into the view of a camera.

4.3 Target Video Retrieval and Social Group Clustering

We addressed target video retrieval and social group clustering as other tasks enabled by our approach. More specifically, given a collection of videos that combined our dataset (scenes **Indoor 0 – Indoor 3** and **Outdoor 0 – Outdoor 3**) and the CMU dataset (scenes **Pool, Tennis, Chat**), we retrieved videos that included target people (retrieval task) and clustered videos into groups based on the interaction scenes (clustering task).

For both of the tasks, we used the affinity introduced in Eq. (6), which corresponded to **Proposed (C)** in the previous experiment in that we relied on the only correlation-based targetness. We also implemented some other forms of affinity that used (1) $P(a_{X^{(i)}} \mid \mathcal{V}_p)$ instead of $P(\mathcal{V}_p \mid a_{X^{(i)}})$ in Eq. (5) corresponding to **Proposed (C+G)** in the previous experiment and (2) an asymmetric affinity $A(\mathcal{V}_p \mid \mathcal{V}_q)$ instead of $A(\mathcal{V}_p, \mathcal{V}_q)$ in Eq. (6), which we refer to as **Proposed (asym)**. We calculated the average length of target candidates across all the videos for μ_l. Because many objects in each scene were different, we implemented an unsupervised scene clustering as a baseline. Specifically, we used the GIST

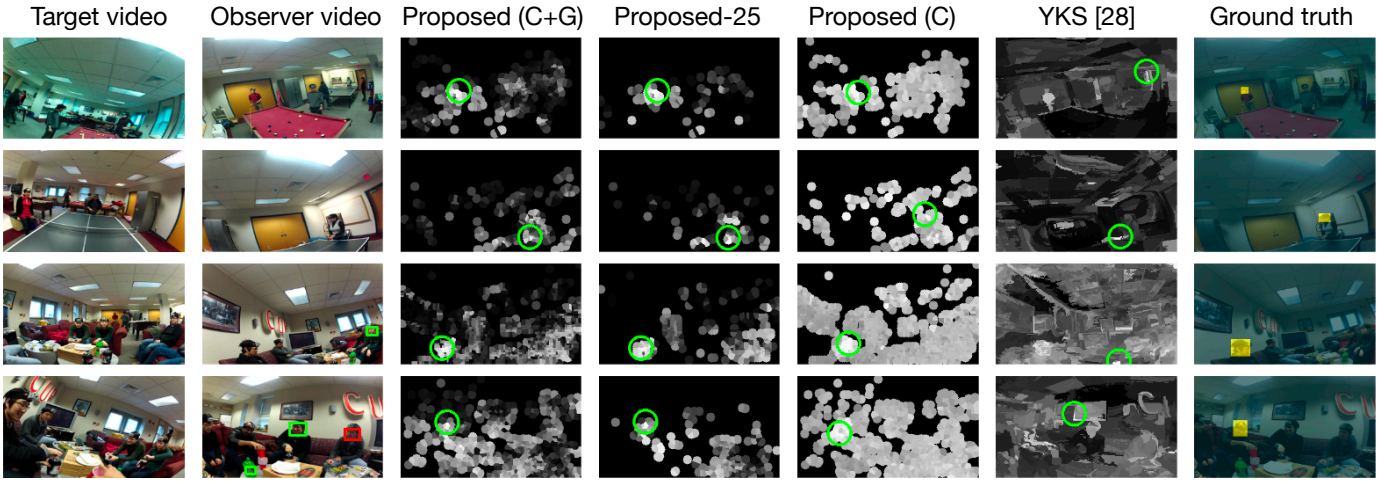

Fig. 5. Target localization results on the CMU dataset. The target candidate with the highest targetness score is specified by the green circle on **Proposed (C+G)**, **Proposed-25**, **Proposed (C)** and **YKS** [28]. **ZR**: Mixtures-of-tree based facial landmark detector [36]. **VJ**: Haar-cascade face detector [37] (green rectangles in the second column). **VJ + LBPH**: face recognition using local binary pattern histograms [38] (red rectangles). **VJ + FisherFace**: face recognition using the Fisher face [39] (yellow rectangles).

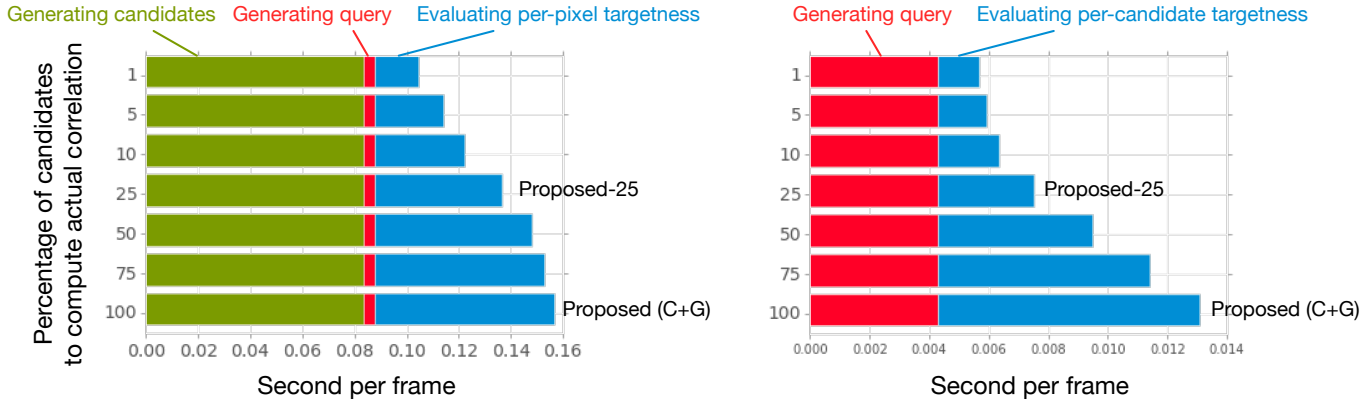

Fig. 7. Comparison of computation times on different percentages of candidates to evaluate an actual correlation. Left: overall computation times. Right: computation times to evaluate per-candidate targetness where the candidates are already generated.

scene descriptor [41][5] and Places CNN [42][6]. These features were extracted every 50 frames and were averaged over time to form a feature vector for each video.

Table 3 demonstrates the results of target video retrieval. As the number of videos including target people was known for each target person, we evaluated the mean R-precision score [43] for each method. Namely, we calculated the precision@R for each target video, where R is the number of videos including the target. Tab. 4 describes the results of social group clustering. As the number of groups was practically unknown, we used the affinity propagation algorithm [44] to estimate both the number of groups and the group assignment for each video. We evaluated the precision, recall and f-measure scores as well as the estimated number of groups.

Overall, we found that **Proposed (C)** performed the best. Fig. 9 depicts the affinity matrices used for the two tasks. As we mentioned in Sec. 2.4, the use of generic targetness in **Proposed (C+G)** made affinities similar to each other. On the other hand, high affinity scores could be found in a block diagonal manner when using **Proposed (C)**, and they were more similar to the ground truth affinity depicted in the right of the figure. We also found the performance degraded in **Proposed (asym)**. This result indicates the importance of targetness of an opposite direction: how likely observers are included in the videos recorded by a target. Note that exactly the same affinity matrices could be obtained in a much shorter time using the two-step search scheme with $P = 25$ (63 percent reduction in computation times as reported in Sec. 4.2).

5 Conclusions

We introduced a novel correlation-based approach to the problem of self-search for first-person videos. Experimental results demonstrate that our proposed method was able to search and localize self instances robustly without the use of

5. We used the code available at http://lear.inrialpes.fr/software.

6. We used the pre-trained model Places205-GoogLeNet available at http://places.csail.mit.edu/downloadCNN.html. Responses in the last layer (*i.e.*, posterior probabilities for each scene category) were used for the features.

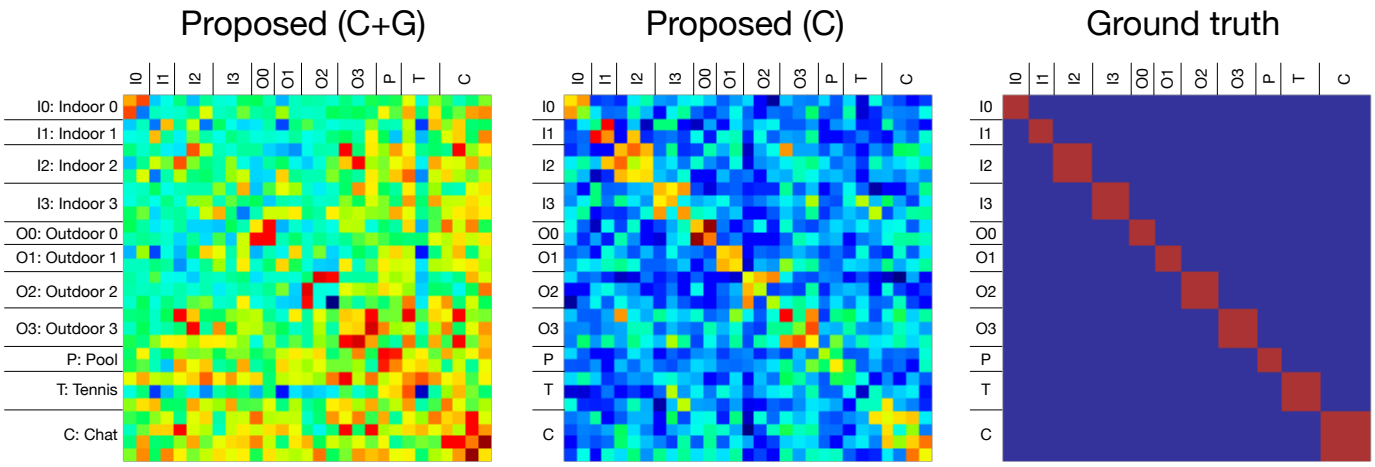

Fig. 9. Affinity matrices between first-person videos used for target video retrieval and social group discovery.

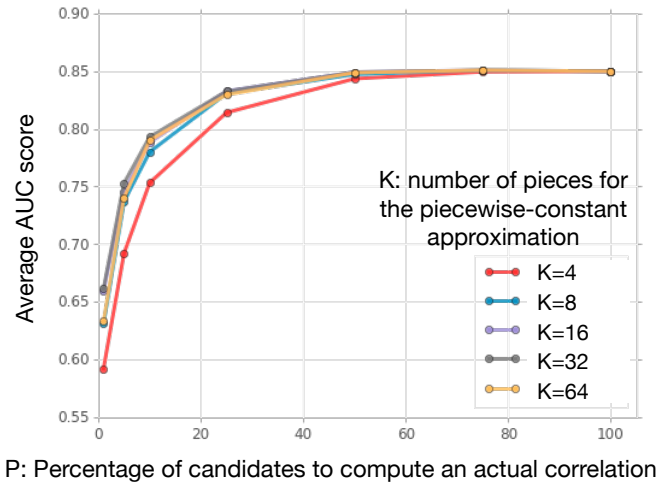

Fig. 6. Comparison of AUC scores on different combinations of P: the percentages of candidates to evaluate an actual correlation and K: the numbers of pieces for the piecewise-constant approximation.

TABLE 3
Comparison of the mean R-precision on target video retrieval.

Proposed (C+G)	Proposed (C)	Proposed (asym)
0.48	**0.61**	0.50
GIST [41]	**Places CNN** [42]	
0.29	0.27	

Fig. 8. Target segmentation based on per-pixel targetness maps to avoid violating the privacy of other people who accidentally come into the view of the target's camera.

face recognition. Trajectory-based target candidates enabled comparably accurate and efficient searches over supervoxel-based candidates presented in our prior work [28]. The upperbound estimation of correlation scores was used to limit the number of candidates to apply a stable but slow evaluation of correlation-based targetness, making an overall search procedure more efficient while keeping high search performance.

One limitation of our approach is that it is not well suited for crowded scenes, where many people may be moving in the same way (*i.e.*, high correlation across multiple individuals) and where individuals are only visible for short periods of time (*i.e.*, not enough signal). Searching for people in these types of videos will require a richer set of features and will be an interesting direction for future work.

TABLE 4
Comparison of clustering performance on social group clustering. Note that the ground-truth number of groups was 11.

	Precision	Recall	F-measure	#groups
Proposed (C+G)	0.47	0.70	0.54	7
Proposed (C)	**0.81**	**0.89**	**0.84**	**10**
Proposed (asym)	0.59	0.65	0.60	**10**
GIST [41]	0.24	**0.89**	0.36	3
Places CNN [42]	0.42	0.74	0.49	6

Extending the self-search to a larger repository of first-person videos will enable many novel applications. For example, searching for a group of people across first-person videos recorded at a variety of places around the world will illuminate their social activities, which have never been pursued by any visual surveillance. This application also raises new computer vision problems, such as social saliency prediction [27] on a wide-spread area and group activity summarization for large-scale first-person videos.

APPENDIX A
IMPLEMENTATION DETAILS

In our experiments, we resized video frames into 320x180 to complete an overall procedure in reasonable time. The spatial step e_W to sample trajectories was set to four pixels. The good-feature-to-track criterion, *i.e.*, the minimum eigenvalue of the autocorrelation matrix, was set to 10^{-4}. Every four frames, new trajectories were sampled at the locations satisfying the aforementioned criteria, and where current trajectories do not already exist. Smaller values for e_W and the good-feature-to-track criterion could allow us to use more densely-sampled trajectories. However, we found that such settings slightly improved performance at a cost of longer computation time. $L_{\min}$ and $L_{\max}$ were set to 64 and 1024 frames, respectively. The average trajectory length (μ_l in Sec. 2.4) was 220.14 frames, larger than the 200 frames that [18] claimed to ensure low false-positive rates.

APPENDIX B
DERIVING AN UPPERBOUND OF CORRELATION

In this appendix, we introduce how to derive Eq. (9) from Eq. (8). By expanding the LHS of Eq. (9), we obtain

$$\frac{1}{l}\sum_t (u_t - U_t)^2 = \frac{1}{l}\sum_t (u_t^2 - 2u_tU_t + U_t^2) \quad (12)$$

$$= 2(1 - \frac{1}{l}\sum_t u_tU_t). \quad (13)$$

In Eq. (13), recall that $\boldsymbol{u}$ and $\boldsymbol{U}$ are normalized to have zero mean and unit variance, and thus $\frac{1}{l}\sum_t u_t^2 = \frac{1}{l}\sum_t(u_t^2 - \text{mean}(\boldsymbol{u})) = \sigma_{\boldsymbol{u}}^2 = 1$, and $\frac{1}{l}\sum_t U_t^2 = 1$. Likewise, we can expand the RHS of Eq. (8) as follows:

$$\frac{1}{K}\sum_k (\bar{u}_k - \bar{U}_k)^2 = \sigma_{\boldsymbol{u}}^2 + \sigma_{\boldsymbol{U}}^2 - \frac{2}{K}\sum_k \bar{u}_k\bar{U}_k. \quad (14)$$

By substituting Eq. (13) and Eq. (14) into Eq. (8), Eq. (9) is derived.

ACKNOWLEDGMENTS

This research was supported by CREST JST and the Kayamori Foundation of Informational Science Advancement.

REFERENCES

[1] I. Arev, H. S. Park, Y. Sheikh, J. Hodgins, and A. Shamir, "Automatic editing of footage from multiple social cameras," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, pp. 81:1–81:11, 2014.
[2] Y. J. Lee, J. Ghosh, and K. Grauman, "Discovering important people and objects for egocentric video summarization," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 1346 – 1353.
[3] Z. Lu and K. Grauman, "Story-driven summarization for egocentric video," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 2714–2721.
[4] J. Xu, L. Mukherjee, Y. Li, J. Warner, J. M. Rehg, and V. Singh, "Gaze-enabled egocentric video summarization via constrained submodular maximization," 2015, pp. 2235–2244.
[5] T.-S. Leung and G. Medioni, "Visual navigation aid for the blind in dynamic environments," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2014, pp. 153 – 158.
[6] T. J. J. Tang and W. H. Li, "An assistive eyewear prototype that interactively converts 3d object locations into spatial audio," in *Proc. ACM International Symposium on Wearable Computers (ISWC)*, 2014, pp. 119–126.
[7] Y. Tian, Y. Liu, and J. Tan, "Wearable navigation system for the blind people in dynamic environments," in *Proc. Cyber Technology in Automation, Control and Intelligent Systems (CYBER)*, 2013, pp. 153 – 158.
[8] E. Hjelmås and B. K. Low, "Face detection: A survey," *Computer Vision and Image Understanding (CVIU)*, vol. 83, no. 3, pp. 236 – 274, 2001.
[9] C. Zhang and Z. Zhang, "A survey of recent advances in face detection," Tech. Rep. MSR-TR-2010-66, 2010.
[10] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face recognition: A literature survey," *ACM Computing Surveys*, vol. 35, no. 4, pp. 399–458, 2003.
[11] S. Gong, M. Cristani, C. Loy, and T. Hospedales, "The re-identification challenge," in *Person Re-Identification*, ser. Advances in Computer Vision and Pattern Recognition. Springer London, 2014, pp. 1–20.
[12] R. Vezzani, D. Baltieri, and R. Cucchiara, "People reidentification in surveillance and forensics: A survey," *ACM Compututing Surveys*, vol. 46, no. 2, pp. 29:1–29:37, 2013.
[13] H. Wang and C. Schmid, "Action recognition with improved trajectories," in *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2013, pp. 3551 – 3558.
[14] P. Salvagnini, L. Bazzani, M. Cristani, and V. Murino, "Person re-identification with a ptz camera: An introductory study," in *Proc. IEEE International Conference on Image Processing (ICIP)*, 2013, pp. 3552 – 3556.
[15] S. Alletto, G. Serra, S. Calderara, and R. Cucchiara, "Head pose estimation in first-person camera views," in *Proc. International Conference on Pattern Recognition (ICPR)*, 2014, pp. 1–6.
[16] S. Alletto, G. Serra, S. Calderara, F. Solera, and R. Cucchiara, "From ego to nos-vision: Detecting social relationships in first-person views," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2014, pp. 594–599.
[17] A. Fathi, J. K. Hodgins, and J. M. Rehg, "Social interactions: A first-person perspective," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 1226–1233.
[18] Y. Poleg, C. Arora, and S. Peleg, "Head motion signatures from egocentric videos," in *Proc. Asian Conference on Computer Vision (ACCV)*, 2014, pp. 1–15.
[19] J. Rehg, G. Abowd, A. Rozga, M. Romero, M. Clements, S. Sclaroff, I. Essa, O. Ousley, Y. Li, C. Kim, H. Rao, J. Kim, L. Presti, J. Zhang, D. Lantsman, J. Bidwell, and Z. Ye, "Decoding children's social behavior," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 3414 – 3421.
[20] Z. Ye, Y. Li, A. Fathi, Y. Han, A. Rozga, G. D. Abowd, and J. M. Rehg, "Detecting eye contact using wearable eye-tracking glasses," in *Proc. ACM Conference on Ubiquitous Computing (UbiComp)*, 2012, pp. 699–704.

[21] Z. Ye, Y. Li, Y. Liu, C. Bridges, A. Rozga, and J. M. Rehg, "Detecting bids for eye contact using a wearable camera," in *Proc. IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, 2015.

[22] S. Alletto, G. Serra, S. Calderara, and R. Cucchiara, "Understanding social relationships in egocentric vision," *Pattern Recognition*, vol. 48, no. 12, pp. 4082–4096, 2015.

[23] J. Hesch and S. Roumeliotis, "Consistency analysis and improvement for single-camera localization," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2012, pp. 15 – 22.

[24] A. Murillo, D. Gutierrez-Gomez, A. Rituerto, L. Puig, and J. Guerrero, "Wearable omnidirectional vision system for personal localization and guidance," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2012, pp. 8 – 14.

[25] H. S. Park, E. Jain, and Y. Sheikh, "3d social saliency from head-mounted cameras," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2012, pp. 1–9.

[26] ——, "Predicting primary gaze behavior using social saliency fields," in *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2013, pp. 3503 – 3510.

[27] H. S. Park and J. Shi, "Social saliency prediction," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 4777–4785.

[28] R. Yonetani, K. M. Kitani, and Y. Sato, "Ego-surfing first-person videos," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 5445–5454.

[29] C. Xu and J. J. Corso, "Evaluation of super-voxel methods for early video processing," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 1202–1209.

[30] C. Xu, C. Xiong, and J. Corso, "Streaming hierarchical video segmentation," in *Proc. European Conference on Computer Vision (ECCV)*, 2012, pp. 1–14.

[31] J. Shi and C. Tomasi, "Good features to track," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1994, pp. 593 – 600.

[32] G. Farnebäck, "Two-frame motion estimation based on polynomial expansion," in *Proc. Scandinavian Conference on Image Analysis (SCIA)*, 2003, pp. 363–370.

[33] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, 1981, pp. 674–679.

[34] C. Faloutsos, M. Ranganathan, and Y. Manolopoulos, "Fast subsequence matching in time-series databases," *SIGMOD Record*, vol. 23, no. 2, pp. 419–429, 1994.

[35] E. Keogh, K. Chakrabarti, M. Pazzani, and S. Mehrotra, "Dimensionality reduction for fast similarity search in large time series databases," *Knowledge and Information Systems*, vol. 3, no. 3, pp. 263–286, 2001.

[36] X. Zhu and D. Ramanan, "Face detection, pose estimation and landmark localization in the wild," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 2879 – 2886.

[37] P. Viola and M. Jones, "Robust real-time object detection," *International Journal of Computer Vision (IJCV)*, vol. 57, no. 2, pp. 137–154, 2004.

[38] T. Ahonen, A. Hadid, and M. Pietikainen, "Face description with local binary patterns: Application to face recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 28, no. 12, pp. 2037–2041, 2006.

[39] P. N. Belhumeur, J. a. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. fisherfaces: Recognition using class specific linear projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 19, no. 7, pp. 711–720, 1997.

[40] C. Rother, V. Kolmogorov, and A. Blake, ""grabcut": Interactive foreground extraction using iterated graph cuts," *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 309–314, 2004.

[41] A. Torralba, K. Murphy, W. Freeman, and M. Rubin, "Context-based vision system for place and object recognition," in *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2003, pp. 273 – 280.

[42] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, "Learning deep features for scene recognition using places database," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2014, pp. 487–495.

[43] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*. Cambridge University Press, 2008, ch. 8: Evaluation in information retrieval.

[44] B. J. Frey and D. Dueck, "Clustering by passing messages between data points," *Science*, vol. 315, no. 5814, pp. 972–976, 2007.

Ryo Yonetani received his MS and PhD degrees in informatics from Kyoto University in 2011 and 2013. He is currently a research associate at Institute of Industrial Science, the University of Tokyo. His research interests include first-person vision, visual attention, and computer-human interaction.

Kris M. Kitani is a Systems Scientist at the Robotics Institute at Carnegie Mellon University. He received his BS degree from the University of Southern California, and MS and PhD degree from the University of Tokyo. His research interests include first-person vision, activity forecasting and visual compliers.

Yoichi Sato is a professor at Institute of Industrial Science, the University of Tokyo. He received his B.S. degree from the University of Tokyo in 1990, and his MS and PhD degrees in robotics from School of Computer Science, Carnegie Mellon University in 1993 and 1997. His research interests include physics-based vision, reflectance analysis, first-person vision, and gaze sensing and analysis. He served/is serving in journal editorial roles including ones for IEEE Transactions on Pattern Analysis and Machine Intelligence, International Journal of Computer Vision, and Computer Vision and Image Understanding.